13

Strategies for Interactive Task Learning and Teaching

Katrien Beuls, Luc Steels, and Paul Van Eecke

Abstract

A strategy is a way to make decisions that come up when handling a task. It requires a problem solver able to address routine cases and a set of diagnostics and repairs to handle, in a flexible way, unusual or unforeseen situations. Between humans, interactive task learning and teaching appear to involve strategies at three levels: (a) the execution of a task with available knowledge (task strategy), (b) interactive learning to expand the available knowledge and thus become a better problem solver in the future (learning strategy), and (c) interactive teaching or tutoring to help others learn (teaching strategy). This chapter examines the general architecture that is needed to build artificial agents that can play either the role of teacher, by carrying out teaching strategies, or the role of learner, by carrying out learning strategies that benefit from these teaching strategies. Focus is on artificial teachers that interact with humans or artificial learners as well as on artificial learners that interact with human or artificial teachers. We argue that the use of a meta-layer is of primary importance for understanding and implementing strategies and point to operational examples from an implementation of this hypothesis in the domain of second-language teaching.

Introduction

Humans invent and acquire an amazing number of day-to-day tasks and routinely execute them throughout life. Some of these tasks, such as medical surgery, manufacturing, and business procedures, may be more professional in purpose but they rely, nevertheless, on the same cognitive capacities as "commonsense" tasks, such as cooking or cleaning a room. Some task execution procedures have been designed or analyzed explicitly; they have been written down and are taught through explicit education. Other procedures, however, are acquired through interactions with people who have the requisite knowledge to carry out the task.

Our discussion here centers on the cognitive process and interaction patterns by which tasks are learned and taught in an interactive-situated manner,

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

at a sufficiently deep level so that we can emulate these processes in artificial systems. The notion of a strategy plays a central role in this endeavor, and the expertise needed for task execution, learning, and teaching relies on domain knowledge and strategies, builds up a context model, and performs goal/subgoal decompositions. We begin by examining conceptual issues and then discuss an application, the *Spanish Verb Tutor*, in an intelligent tutoring system that is designed to assist second-language learning.

What Are the Components of Task Execution and Task Learning?

It is possible, and indeed common in artificial intelligence, to view learning as a particular kind of problem solving (Mitchell et al. 1983; Simon 1996). There is often more than one way to generalize or specialize a concept or inference and often more than one possible extension of a strategy or different hypotheses about domain models. Just as nontrivial tasks require consideration of different possible avenues to tackle a task and the heuristics to make a decision, these different options must be evaluated based on learning heuristics to derive the most likely option to pursue. Given this perspective, we can ask whether the same components found in the kind of problem solving involved in the execution of tasks are also found in learning (Steels 1990).

Task Expertise

From diagnosing failure in a mechanical system, to writing a computer program, or deciding whether an insurance claim is valid and assessing how much needs to be paid, a task requires coming up with a series of actions to address the peculiarities of that task in a given context. These actions must satisfy the *goals* and *subgoals* required by the task.

Some tasks are so routine that there is a ready-made solution at hand. Here, however, we are interested in tasks that require problem solving. This involves obviously a fair amount of facts and inference rules about the domain (captured in a *domain model*) and facts about the particular case in which the task is to be executed (captured in a *context model*). These facts are formulated in an ontology that defines the categories with which to describe the objects and properties that make up these models. In addition, problem solving requires a *task strategy* for making use of the domain model to expand the context model and the goal/subgoal decomposition to achieve the task.

We refer to the set of domain ontologies, domain models, and task strategies as the *domain knowledge* of a particular task domain. The contents of domain knowledge are relevant across different tasks in the same domain, whereas the context model and goal/subgoal decomposition are specific to a concrete task context. Elements of domain knowledge typically have several degrees of generality. The ontology and domain models consist partly of concepts and

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

facts of which a great deal is valid across domains, whereas some of it is more specialized. Similarly task strategies range from very specific strategies, which work to solve only a limited set of tasks in a very specific contextual setting, to more generic approaches, such as "divide and conquer" or "simplify and extend" (Polya 1945).

It is common to add an additional component to a task execution agent, a *meta-layer*, so that there are two levels of problem solving (Figure 13.1):

1. Basic processing where available domain knowledge is applied and possibly leads to acceptable performance.
2. Meta-level processing, which becomes active when basic processing fails, in which case the problem solver moves to a higher level to apply the available knowledge much more flexibly, for instance, by relaxing certain constraints so that an inference schema that does not fit completely with the current situation can still be applied, or by ignoring gaps in the input and continue to find the best possible solution.

These strategies constitute metaknowledge about the task domain, and architectures to represent and operationalize these strategies have been studied intensely since the 1980s (e.g., Rosenbloom et al. 1986). Meta-level strategies can also invoke learning expertise to expand the domain knowledge and consolidate adaptations. The different components of task expertise are summarized in Table 13.1 (Steels 1990).

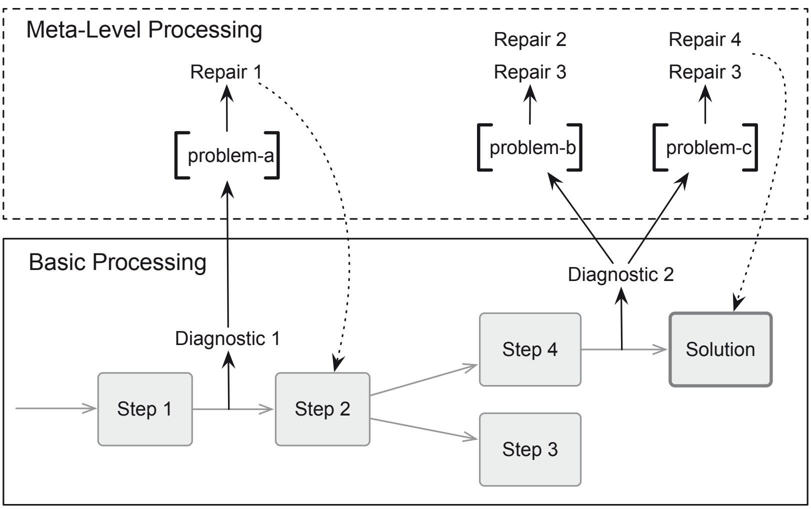

Figure 13.1 Diagnostics and repairs operate on top of basic problem solving, which traverses a search space of possible steps toward a solution. Diagnostics signal problems (e.g., problem-a, problem-b). A diagnostic can be handled by a range of repair strategies. The repair strategies are compared, the most plausible strategy is chosen for execution, and after the repair the basic layer is evoked again to continue or restart from an earlier step in problem solving.

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

Table 13.1 Components of expertise, found both in task expertise and learning expertise.

Task General	Task Specific
Ontologies	
Domain models	Context models
Strategies	Goal/subgoal decomposition

Learning Expertise

Can we find the analogues of goals, ontologies, domain models, strategies, context models, and goal/subgoal decomposition in learning?

The expansion of knowledge in a particular task domain (the domain ontologies, domain models, and task strategies) is the *learning goal*. For example, when an apprentice receives instruction to become capable of operating a complex machine, this learning task becomes an explicit goal, and the learner will need to decompose this goal into many learning subgoals. Of course, a lot of learning also takes place in the background, triggered when there are opportunities for generalization or specialization or when new domain facts arise. In this case, we would say that these learning goals are continuously active.

There are many *learning strategies* and among them there are important individual differences. A learning strategy is a series of steps that are followed to acquire some aspect of task expertise. For example, the learner could chunk some of the inference steps into a more compact single rule to speed up problem solving later, and thus handle more complex problems; alternatively, the learner could detect similarities between a set of objects and introduce a category to be able to represent inferences that apply to the class of objects of the same category or seek additional information from external sources. Many of these learning strategies have been operationalized in the symbolic machine learning literature and classified in terms of dimensions such as the degree of supervision (unsupervised/supervised), the learning objectives (similarity, classification, discrimination), the nature of the domain model (decision trees, weighted networks, inferences rules), the nature of available data, and the availability of automatic mechanisms cooperating with deliberate strategies (Mitchell 1997).

The *learning ontology* provides the building blocks for formulating *learning models* that capture general facts and experiences about learning, in particular which strategies are most effective for certain learning tasks. For example, the acquisition of part of a particular domain ontology might be based on decision-tree learning, whereas the acquisition of heuristics might be based on exploring a search space and storing information on major decision points based on their success in a task; deblocking an impasse in a specific task domain might be effective by relaxing constraints so that inference rules could apply more flexibly, whereas in another domain it could be by relaxing some supposed

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

constraints on the task itself. The ontology is also required for building the rich *learning context models* needed in learning. These learning context models represent information about a specific learning episode so that the learner can make decisions on how to handle it.

How Do Task Execution and Task Learning Interact?

Much recent research in machine learning, particularly in statistical and neural network learning, assumes that there is a strict separation between a learning phase, based on cycling through a large amount of data, and an execution phase (Bengio 2009). However, this is certainly not the only form of learning. Interactive task learning is situated within a concrete context. It requires incremental learning strategies, which absorb information from a concrete case and are implemented in tight interaction with the problem-solving processes required for task execution. How can such a tight interaction be achieved?

A plausible solution is to assume that the *meta-layer architecture*, introduced earlier for implementing more flexible problem solving, can also be used for invoking learning strategies. At the meta-level, the agents reason about gaps in the domain knowledge and how they can be filled by invoking an appropriate learning strategy. After the learning strategy has taken its course, processing can move back to the task execution layer and the newly available knowledge can then be used to pursue problem solving and task execution further.

The move back to the meta-level gets triggered when a learning opportunity arises. This can occur, for instance, when there is an error in performance, signaled through feedback from the environment (e.g., an action did not generate the expected consequences), when there are impasses during the execution of a task (e.g., the learner gets stuck trying to achieve a particular domain goal), or when an opportunity is sensed for integrating new knowledge into the existing body of domain knowledge (e.g., by generalization or specialization after handling a concrete case).

How Does Teaching Expertise Relate to Task Learning and Execution?

Although learning and teaching are often seen as two different activities that can or should be studied separately, we argue that it is better to view them as two sides of the same coin, just as language understanding and language production are intimately intertwined and the same competence intervenes for both. Those who are better learners tend to be better teachers because they have more successfully developed metacognition skills, such as diagnosing and repairing knowledge gaps (Veenman et al. 2006). Conversely, learners can become better by acquiring skills associated with teaching, such as

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

setting learning goals, timing, and motivation (Zimmerman 2010). We argue that learning and teaching share general characteristics with problem solving, which implies that the same fundamental components (as discussed above) can be expected. Indeed, in the educational literature, teaching expertise has already been viewed as a problem-solving/decision-making process (Shuell 1990), and there have been extensive efforts to document teaching strategies and study their effectiveness in many domains. Moreover, there are considerable differences in the strategies that different individuals use to teach, and it is possible to become a better teacher through practice and instruction. Thus, teaching expertise can be approached in the same way as task expertise (Steels and Tokoro 2003).

Some teaching strategies involve the systematic presentation of domain knowledge and the introduction of ways to exercise and test whether this knowledge has been acquired by the learner (see VanLehn, this volume). Traditional top-down teaching strategies that dominate classroom teaching exemplify these types of strategies. Other teaching strategies employ a more active, learner-oriented approach. They assume that the teacher monitors carefully what the learner already knows, or what kind of errors the learner is making, and intervenes with examples, challenges, and corrections that zoom in on these issues. The learner-oriented approach is a more natural form of teaching and is practiced by parents, caregivers, or peers in natural learning settings. It is particularly this form of interactive teaching that interests us, even though it is much more difficult to capture in computational learning environments. More sophisticated teaching strategies use a model of what the learner already knows to plan possible exercises or gauge whether new material can already be presented (Amaral and Meurers 2007; VanLehn 1988).

The implementation of teaching strategies requires the same components as for task execution and learning: an ontology to describe teaching situations, domain models that now focus on capturing knowledge about teaching, a context model that represents the current teaching episode, teaching goals and subgoals, and teaching strategies. Teaching expertise can either be in the driver seat of an interaction (e.g., in a classroom situation where a teacher has explicit teaching goals and then uses strategies to present new material or come up with appropriate exercises) or invoked when discrepancies are discovered between the behavior of the learner and the behavior that the teacher was expecting.

Application in Intelligent Tutoring Systems for Second-Language Learning

The insights and proposals in this chapter derive partly from our technical work on language speaking, understanding, learning, and teaching (Beuls 2014; Steels and Hild 2012). Speaking and understanding language can be seen as very complex tasks that have all the characteristics of other commonsense

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

tasks. Speaking is similar to a planning or design task: it requires that a (communicative) goal is decomposed into various subgoals and the speaker needs to find the best way to translate some aspect of meaning into words and syntactic structures that conform to the conventions of the language. Understanding is similar to a plan recognition task: the listener must grasp the purpose of each word or syntactic structure and interpret it within the present context.

Young children, being native language learners, rely heavily on experience; they memorize stereotyped patterns and only gradually systematize them. They do not get explicit instruction, and their knowledge of the underlying structure of their native language remains implicit and cannot be verbalized (e.g., Dabrowska and Lieven 2005). Second-language learners, by contrast, are given instructions on the verbal paradigms, phrase structures, and many other aspects of the language that they are trying to learn, and they then gradually internalize these rules so that they become routine (cf. multiple contributions in Robinson and Ellis 2008). They also need a lot of practice and experience, which seems contradictory to the pedagogical aim of second-language learning: learning is to be achieved at an accelerated pace with less input than natural learning, even though this is not always successful. Language is a good test bed to explore experiential and symbolic knowledge acquisition and ways in which they interact.

In this context, we briefly introduce a concrete example of a second-language teaching application designed for learning Spanish verbs (for details, see Beuls 2013). This application, the *Spanish Verb Tutor*, illustrates the three strategy levels (task execution, learning, and teaching) as well as the use of meta-level processing.[1]

The task for the human learner is to engage in language games that approach, as much as possible, real-world situations in which verb conjugation is relevant. Two events are visually shown on a time line containing a past, present, and future time sphere. The time line may also contain the actor in the event, described as the subject of the verb (see Figure 13.2). In a game, the human learner goes through a series of situations and has to act either as a listener or as a speaker.

When playing the listener role, the learner gets a sentence with a conjugated Spanish verb and the task consists of choosing which event is in the correct location on the time line (see Figure 13.2a). As a speaker, the learner selects an event on the time line and then enters the correct verb form to meet this temporal condition (see Figure 13.2b). If the required skill level is too high, the learner can choose to skip the lesson. After each interaction, the learner receives feedback on the answer and an explanation of the type of error.

The *Spanish Verb Tutor* is implemented using two agents (see Figure 13.3): a Teacher Agent (in charge of teaching) and a Learner Agent (in charge of

[1] For a demonstration of the system, see https://www.youtube.com/watch?v=5BWDVGUjEEs (accessed January 23, 2019).

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

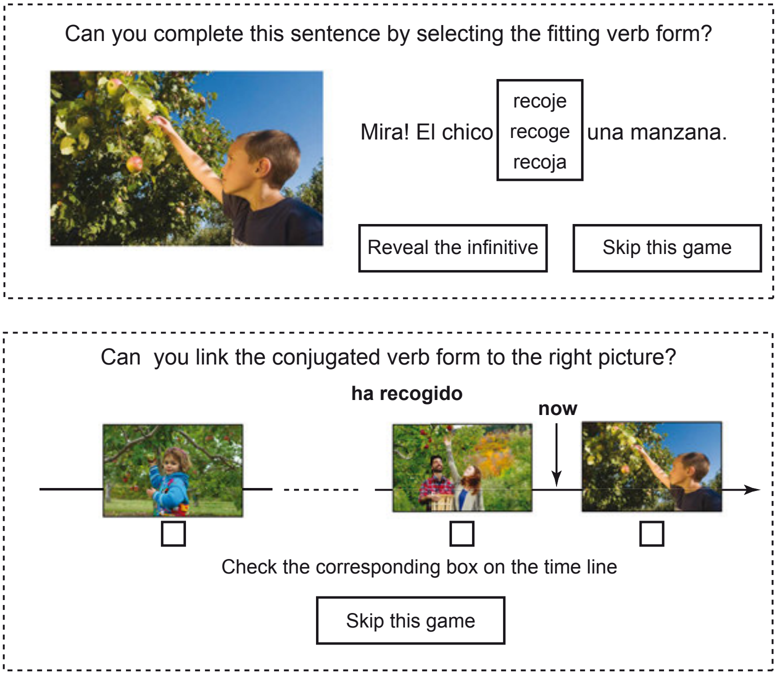

Figure 13.2 Screen shots of the *Spanish Verb Tutor*, illustrating (a) comprehension and (b) production language tasks.

learning a model of the human learner). The Learner Agent is used by the Teacher Agent to personalize the teaching experience of the human learner. Each agent has two components.

The first component involves task execution. Both the Teacher Agent and the Learner Agent have *active expertise* in the sense that they are able to play the language game autonomously. The Teacher Agent, however, plays the game equipped with a fully competent, accurate understanding of the grammar of the *Spanish Verb* system, whereas the Learner Agent models a human learner's competence in Spanish verb conjugation, which is necessarily partial and erroneous. The language behavior of the agents is implemented using fluid construction grammar (FCG) (Steels 2017), but it could be implemented in any other formalism with similar functionalities. To play the game, *flexibility* is required from both agents. The Teacher Agent will be confronted with ungrammatical or partial input from the human learner and requires flexibility to respond properly. The Learner Agent needs flexibility to parse sentences produced by the Teacher Agent, when the learner's own grammar (being a model of the grammar of the human learner) is unable to answer correctly. Flexibility is implemented using the meta-level functionalities embedded in the FCG formalism (Van Eecke and Beuls 2017), which uses operators that return a description of what constraints had to be relaxed.

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

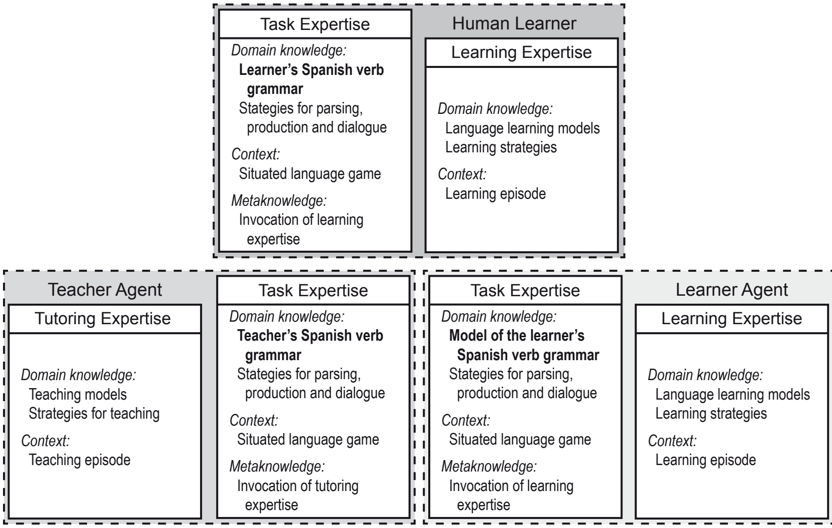

Figure 13.3 The main architecture of the *Spanish Verb Tutor* consists of two agents that interact: the Learner Agent, which is an active model of the human learner, and the Teacher Agent, which can take on the role of teacher. The Learner Agent has task expertise to predict how the human learner will perform in a game, and it has learning expertise to learn the grammar of the human learner to become better at predicting his behavior. The Teacher Agent has tutoring expertise as well as task expertise to formulate exercises or correct human behavior in the task. Teaching strategies either get invoked in a top-down manner or through meta-level operators.

The second component permits the Learner Agent and Teacher Agent to implement learner expertise and task expertise, respectively. Learner expertise operates on a meta-level compared to the task execution level.

Recall that the goal of the Learner Agent is to approximate, to the greatest extent possible, what the human learner already knows, so that teaching strategies can make use of this model to personalize feedback, pinpoint new material, and formulate the most productive exercises. When the Teacher Agent launches a new language game interaction, the Learner Agent carries out the interaction in parallel with the human learner and then compares its response with the human response. If the response is different, the Learner Agent will use learning strategies to implement its model of the human learner, including, and in particular, the acquisition of the user's error constructions.

The goal of the Teacher Agent is to drive the interaction with the human learner. Based on previous interactions, the Teacher Agent invokes a new language game and uses its task expertise to come up with a solution. So, the exercises given to the learner are not ready-made, and hence boringly predictable for the learner: they are actively constructed, taking into account what the human learner knows. When the response of the human learner does not fit

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

with the teacher's own response, the teacher invokes a teaching strategy that attempts to address the mismatch, explains the relevant grammatical knowledge, and then provides new exercises.

The *Spanish Verb Tutor* demonstrates how a meta-layer architecture can be used to create a second-language tutoring system, one that is fully personalized and capable of teaching language through situated interactions that approach normal use of verb phrases; namely, coming up with descriptions of situations or identifying correctly where an event is located on the time line.

The performance of the Teacher Agent was evaluated on the lowest and highest L2 learner level in the Spanish Learner Language Oral Corpus (SPLLOC). The Teacher Agent (FCG) achieves an accuracy of 58% (percentage of isolated corrected forms that equal the gold standard correction). Our system outperforms the standard MS Word grammar checker by almost 30%. The evaluation was run on individual word forms that were not embedded in a sentence (see Beuls 2014).

Conclusions

In this chapter, we have provided strong analogies between the expertise needed for task execution, learning, and teaching. Each of these relies on domain knowledge and strategies, builds up a context model, and performs goal/subgoal decompositions. We also emphasized the role of meta-level operations in the form of diagnostics that detect needs and opportunities and repairs that handle them. In our practice in building teaching agents in the domain of second-language learning, we have found that meta-level operators can provide a smooth interaction between task execution and learning as well as learning and teaching.

Acknowledgments

This paper was written to provide background for discussions at the Ernst Strüngmann Forum on Interactive Task Learning. We are indebted to Kevin Gluck and John Laird, the conveners of this meeting, and to Julia Lupp, director of the Forum, for creating a most stimulating context to reflect on the issues discussed here. We also thank the many participants of the Forum for their feedback and fascinating discussions. For the writing of this paper, Katrien Beuls and Paul Van Eecke were funded by the Chist-Era ATLANTIS project and Luc Steels by a fellowship from the Institute for Advanced Studies in Catalunya (ICREA) as well as support from the Institute for Evolutionary Biology in Barcelona. The development of the *Spanish Verb Tutor* was funded by the Flemish government agency for Innovation by Science and Technology (IWT from 2009 to 2013).

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.